Consumer Bankruptcy: A Fresh Start

By Igor Livshits, James MacGee, and Michèle Tertilt*

There has been considerable public debate on the relative merits of alternative consumer bankruptcy rules. In the United States, this debate has led to legislation making it more costly for households to declare bankruptcy.[1] In contrast, some European countries have recently relaxed legislation prohibiting the discharge of consumer debt (which meant that unlucky debtors remained liable for past obligations indefinitely) and now allow a partial discharge of debt under restrictive conditions (Johanna Niemi-Kiesilainen 1997; Michelle Alexopoulos and Ian Domowitz 1998).

This paper contributes to this debate by quantitatively analyzing two different consumer bankruptcy arrangements. The first system captures key features of Chapter 7 of the US bankruptcy code. We refer to this as a "Fresh Start" (FS) system, since debtors can discharge their debt via bankruptcy and continue their lives free of their existing debt. The second system, which we term "No Fresh Start" (NFS), is motivated by continental Europe. In this system, consumer bankruptcy restructures a consumer's debt payments and limits the amount of earnings that can be garnished.

The quantitative evaluation of consumer bankruptcy laws involves an assessment of the magnitude of two opposing forces. On the one hand, bankruptcy weakens agents' ability to commit to future debt repayment, which limits their ability to smooth consumption across time. On the other hand, in incomplete markets environments, bankruptcy increases households' ability to smooth across states as it introduces some contingency into debt contracts. The easier it is for consumers to discharge their debt, the greater the insurance against "bad luck" such as divorce, job loss, or medical problems. Thus, consumer bankruptcy laws can help consumers smooth their consumption across states at the cost of distorting their ability to smooth over time (see William R. Zame 1993, or Pradeep Dubey, John Geanakoplos, and Martin Shubik 2005).[2] This trade-off implies that any evaluation of bankruptcy regimes must consider the quantitative costs of borrowing constraints versus the value of (partial) insurance against "bad luck."

We undertake our quantitative analysis using a heterogeneous agent life cycle model. Households face both income shocks and expense uncertainty (e.g., uninsured medical bills, divorce costs, or unplanned children).[3] Each period, households make a consumption-savings

* Livshits: Department of Economics, University of Western Ontario, Social Science Centre, London, Ontario, N6A 5C2 (e-mail: livshits@uwo.ca.); MacGee: Department of Economics, University of Western Ontario, Social Science Centre, London, Ontario, N6A 5C2 (e-mail: jmacgee@uwo.ca.); Tertilt: Department of Economics, Stanford University, 579 Serra Mall, Stanford, CA 94305-6072 (e-mail: tertilt@stanford.edu). We thank Michelle Alexoupoulos, V. V. Chari, Russ Cooper, Dean Corbea, Jeremy Greenwood, Larry Jones, Tim Kehoe, Narayana Kocherlakota, Ellen McGrattan, Ed Prescott, Víctor Ríos-Rull, Martin Gervais, and seminar participants at the University of Minnesota, University of Toronto, Simon Fraser University, Pennsylvania State University, Southern Methodist University, Guelph University, York University, Arizona State University, University of Southern California, the Federal Reserve Bank of Minneapolis, the Federal Reserve Bank of Atlanta, Midwest Macro Meetings, Canadian Macro Study Group, the Vienna Macro Workshop, and the SED Meetings for helpful comments. Comments from two anonymous referees and the coeditor are much appreciated. We are also grateful to Jonathan Parker and Pierre-Olivier Gourinchas for supplying us with their estimated life-cycle earnings profile, and to the Federal Reserve Bank of Minneapolis for their support during the writing of this paper. Financial support was provided by the Arts, Humanities and Social Sciences Fund at UWO (Livshits, MacGee), the Social Sciences and Humanities Research Council of Canada (Livshits, MacGee), the National Science Foundation under grant SES 0519324 (Tertilt), and the Stanford Institute for Economic Policy Research (Tertilt). The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

[1] See American Bankruptcy Institute (2005) for the main changes in the "Bankruptcy Abuse Prevention and Consumer Prevention Act of 2005."

[2] A related literature has focused on the implications of economies with complete contingent claims markets and limited enforcement (Timothy Kehoe and David Levine 1993; Narayana R. Kocherlakota 1996).

[3] These shocks are frequently cited by bankrupts as the cause of their bankruptcy.

decision and decide whether or not to file for bankruptcy, taking the bankruptcy rule as given. A bankruptcy rule specifies the amount that can be garnished from households that default, whether discharge of debt is granted, and the "waiting period" before a second bankruptcy is possible. Households can borrow (and save) via one-period noncontingent bonds with perfectly competitive financial intermediaries. Intermediaries are able to observe a household's current income, current level of borrowing, and age. An equilibrium result is that the price of debtors' bonds varies with their current income, age, and level of borrowing. In this paper, we abstract from durable goods and focus on the market for *unsecured* consumer credit.[4]

The quantitative model replicates the (age-specific) bankruptcy filing rates and the debt-to-earnings ratio in the US economy fairly well. We conduct a variety of experiments to assess whether an FS or NFS bankruptcy system is more desirable. Our findings suggest that, for reasonable parameter values, the FS system may indeed achieve higher welfare in the United States. This conclusion is sensitive, however, to both the nature and extent of uncertainty, as well as the life-cycle profile of earnings and family size. A key finding is that if we ignore expense shocks, then a bankruptcy arrangement that severely limits the discharge of debt is better than a US FS system.[5] We also find that larger transitory shocks to income make NFS preferable relative to FS. The more persistent are income shocks, however, the more attractive is FS. The welfare comparison is nonmonotone in the variance of the persistent income shock.

We also find that the life-cycle profiles of earnings and family size matter. The consumption profile is steeper in the FS system, while the variance of consumption is smaller under FS for most age groups. This confirms the intuition that FS facilitates insurance across states, while NFS makes life-cycle smoothing easier. As a result, "flatter" income profiles make FS more attractive than NFS.

Despite the extensive policy debates on the merits of different bankruptcy laws, relatively little work has been done to quantify the effects of alternative consumer bankruptcy provisions. Kartik Athreya (2002) builds on S. Rao Aiyagari (1994) and quantitatively analyzes the effects of bankruptcy laws in an incomplete market exchange economy. He finds that eliminating consumer bankruptcy improves welfare. Wenli Li and Pierre-Daniel Sarte (2006) introduce production and a partially exempt asset into this framework. In contrast to Athreya (2002), they find that eliminating bankruptcy reduces welfare, although introducing means testing would lead to small welfare gains.

In addition to expense shocks, a crucial difference between these papers and ours is the modeling of bond prices. Athreya (2002) and Li and Sarte (2006) assume that all agents can borrow at the same interest rate. This implies that intermediaries could make positive profits by deviating from the equilibrium allocation. To avoid this, we allow interest rates to depend on an agent's type and the amount borrowed.

Satyajitt Chatterjee et al. (2005) also allow interest rates to vary with borrowers' characteristics, and find that introducing means testing into the FS system would lead to welfare gains.[6] The main differences between their paper and ours are that we use a life-cycle model; directly parameterize the expense shocks by looking at data on uninsured medical expenses, divorce, and unexpected children; and consider persistent as well as transitory shocks to earnings. Our analysis shows that these features are important.

The paper is organized as follows. Some background on bankruptcy in the United States is given in Section I. Section II describes the model. The benchmark parameterization and results are presented in Section III. In Section IV we explore the importance of various types of uncertainty. Section V concludes.

[4] A study cited by the National Bankruptcy Review Commission (1997, 136) found that only 5 percent of Chapter 7 cases yielded assets that could be liquidated to repay creditors. This suggests that abstracting from durable goods is reasonable given our focus on Chapter 7 bankruptcy.

[5] One caveat is that we take the size of expense shocks as independent of the bankruptcy system. However, the magnitude of expense shocks may vary with the bankruptcy system, as households decisions to purchase insurance may depend upon the ease with which they can discharge debt.

[6] Marina Pavan (2005) incorporates durables and examines the effect of exemptions on bankruptcy filings.

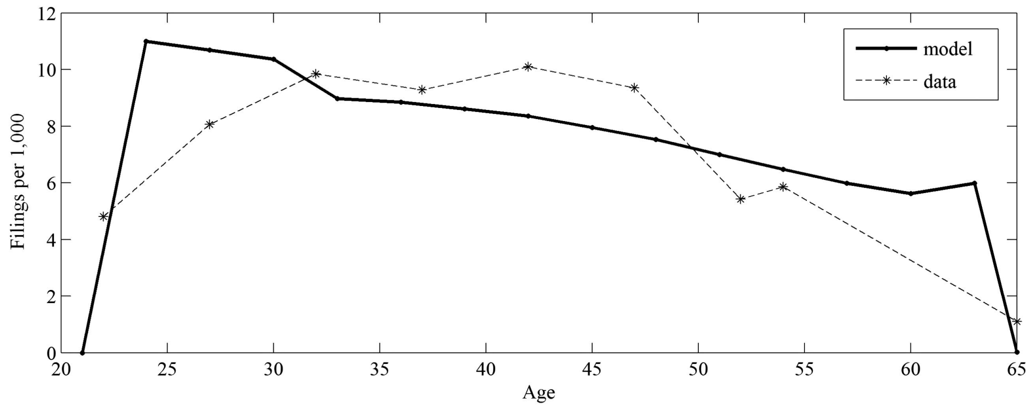

A. Bankruptcies over the Life Cycle

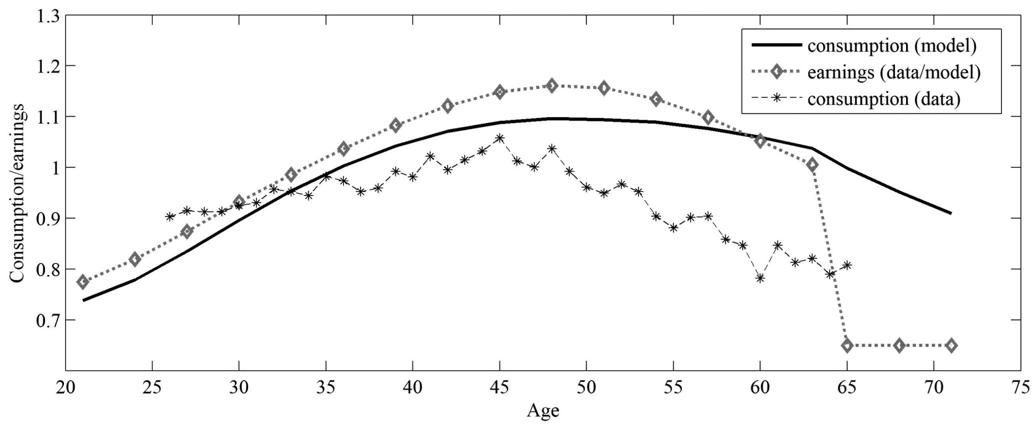

B. Life-Cycle Consumption and Earnings Profiles

FIGURE 1

I. Consumer Bankruptcy in the United States

American households can choose between two bankruptcy procedures: Chapter 7 and Chapter 13.[7] Approximately 70 percent of consumer bankruptcies are filed under Chapter 7. Under Chapter 7, all unsecured debt is discharged in exchange for noncollateralized assets above an exemption level. Debtors are not obliged, however, to use future income to repay debts. Debtors must wait at least six years between Chapter 7 filings. Filers must pay the bankruptcy court filing fee (roughly $200) and legal fees which typically range from $750 to $1,500 (Teresa A. Sullivan, Elizabeth Warren, and Jay Warren Westbrook 2000).

On average, bankrupt households are 30 to 50 percent poorer than the average household, with debt-to-income ratios well above average. Filings first increase, then decrease with age (see Figure 1). The main reported cause of bankruptcy is shocks to income and expenses. Sullivan, Warren, and Westbrook (2000) report that 67.5 percent of filers claimed the main cause of their bankruptcy to be job loss, while 22.1 percent cited family issues such as divorce and 19.3 percent blamed medical expenses (multiple responses were permitted). Other studies find an even larger role for medical

[7] Chapter 13 permits debtors to keep their assets in exchange for a promise to repay part of their debt over three to five years. To qualify for Chapter 13, debtors must have a regular income and secured (unsecured) debts of less than $807,000 ($270,000).

expenses. Elizabeth Warren, Sullivan, and Melissa B. Jacoby (2000) find that 46 percent of filers report either a medical reason or substantial medical debt, while Domowitz and Robert L. Sartain (1999) conclude that medical debt accounts for roughly 30 percent of filings.[8]

II. The Model

We consider an overlapping-generations model where households live for J periods. Each generation is comprised of a continuum of households of measure 1. All households are ex ante identical. They maximize discounted lifetime utility from consumption. Households face idiosyncratic uncertainty, but there is no aggregate uncertainty. Markets are incomplete: the only assets in this economy are person-specific, one-period noncontingent bonds. There are no markets for insurance, and the risk-free interest rate is exogenously given. A crucial element of the model is the household's option to declare bankruptcy.

A. *Households*

Households consume a single good in each period. The preferences are represented by

$$\sum_{j=1}^{J} \beta^{j-1} u\left(\frac{c_j}{n_j}\right), \tag{1}$$

where β is the discount factor, c_j is the total consumption, and n_j is the size of a household of age j in equivalence scale units.[9] We assume that $u(\cdot)$ is strictly increasing and concave.

The labor income of household i at age j, y_j^i, depends upon its productivity and labor endowment:

$$y_j^i = a_j^i \bar{e}_j; \tag{2}$$

$$a_j^i = z_j^i \eta_j^i,$$

where a_j^i is the household's stochastic productivity and $\bar{e}_j$ is the deterministic endowment of efficiency units of labor. The household's productivity is the product of a persistent shock z_j^i and a transitory shock η_j^i. The persistent component z is modeled as a finite Markov chain with an age-independent transition matrix $\Pi(z'|z)$. The productivity of an age 1 household is drawn from the stationary distribution. The transitory component η also has finite support and is i.i.d. over time.

Households face a second type of uncertainty: they may be hit with an idiosyncratic expense shock $\kappa \geq 0$, $\kappa \in K$, where K is the finite set of all possible expense shocks. The probability of shock κ_i is denoted π_i. An expense shock directly changes the net asset position of a household. Expense shocks are i.i.d. and are independent of income shocks.[10]

B. *Financial Markets*

We assume that the risk-free savings interest rate r^s is given exogenously. Loans take the form of one-period bond contracts. The market for bonds is perfectly competitive. The face value of these loans is denoted by d. Note that d is the amount that is to be repaid, not the amount received today. We use the convention that $d > 0$ denotes borrowing, and $d < 0$ denotes savings. Loans are noncontingent as the face value of the loan is not contingent on the realization of any variable. However, the bankruptcy/default option introduces a partial contingency because households have the option

[8] Scott Fay, Erik Hurst, and Michelle White (2002) use the Panel Study of Income Dynamics (PSID) to see how household bankruptcy decisions depend upon household debt, income, and assets. They conclude that "bad luck" is not important, since adding measures of health problems, unemployment, and divorce does not significantly change their results. However, "bad luck" may have a direct impact on debts, income, and assets (Charles Luckett 2002).

[9] The importance of changing family size profile in explaining the hump-shaped life-cycle consumption profile is widely recognized (for example, see Orazio Attanasio and Guglielmo Weber 1995).

[10] This assumption significantly reduces our computational burden. There is also some evidence that suggests that assuming income and expense shocks are independent is not unreasonable. Daniel Feenberg and Jonathan Skinner (1994) (using data from tax returns) find a very low income elasticity of catastrophic health care expenditures, suggesting that expenditure (at least for large medical shocks) does not vary much with income. While one might expect income to decrease in response to a medical shock (by reducing a household's ability to work), our calculations using data from the Medical Expenditure Panel Survey 1996/97 suggest that this effect is small. In particular, we find that average income for households with a large expense shock in 1997, but no shock in 1996, does not decrease. We suspect that this is due to a large fraction of the illnesses being accounted for by children and dependent elderly rather than main wage earners.

of lowering the face value of their debt via bankruptcy.

When making loans, intermediaries observe the total level of borrowing, the current persistent earnings state, and the borrower's age. Thus, the interest rate for borrowers can depend upon age, debt level, and current persistent earnings state. Let $q^b(d, z, j)$ denote the price of a bond issued by a household of age j, with current productivity shock z, and debt d.

Intermediaries maximize expected profits every period. They incur a transaction cost τ of making loans, which is proportional to the size of the loan. In equilibrium, perfect competition assures that intermediaries earn zero expected profits on all loans. This implies that the expected value of repayments must equal the cost of the loan to the intermediary. Perfect competition also implies that, in equilibrium, cross subsidization of interest rates across different types of borrowers will not occur.

C. *Bankruptcy*

A household can declare bankruptcy. A bankruptcy system is characterized by:

- A law of motion for the bankrupt household's debt;
- A repayment rule that specifies the amount of a household's assets and earnings that can be seized by creditors;
- Limited access to financial markets: bankrupts cannot save or borrow during the bankruptcy period.[11]

Our definition of a bankruptcy system incorporates two costs that are frequently mentioned in the literature. One is temporary exclusion from credit markets. In our model, this corresponds to the inability to borrow and save during the bankruptcy period. However, we do not exclude agents from the credit market for any further periods. Although bankruptcy shows up on a consumer's credit report for ten years, many banks specialize in lending to former bankrupts. The second cost is that part of the consumer's income may be seized when bankruptcy is declared. We consider linear garnishment of earnings during the bankruptcy period. The total amount garnished and transferred to creditors is $\Gamma = \gamma y$, where y is earnings and $\gamma \in [0, 1]$ is the marginal rate of garnishment.

We consider two bankruptcy systems. As discussed above, the first, the Fresh Start (FS) system, specifies full discharge of all debts. That is, there is no seizure of future income after the period that bankruptcy is declared. This bankruptcy system captures the key feature of Chapter 7. As in Chapter 7, we do not allow the households to declare bankruptcy more than once within six years. Since households may receive large expense shocks after filing, we assume that households that have large debts but are ineligible to file can "default" but will not have their debt discharged until six years have passed since their last filing. In this case, the household's debt is rolled over for one period and they are subject to the garnishment technology both in the "default" period and in the ensuing period if they file for bankruptcy. Under FS, garnishment is intended to capture the "good faith" requirement of the US bankruptcy code, which we interpret as requiring a certain time of repayment before a borrower can file for bankruptcy.

The second system, No Fresh Start (NFS), is motivated by European bankruptcy laws.[12] The NFS system captures the idea of life-long liability for debt, a key feature of the traditional bankruptcy laws (or lack thereof) in Europe.[13] In this regime, there is no discharge of debt. Instead, a bankrupt's outstanding debt is rolled over at a specified rate of interest $\bar{r}$. This system resembles a repayment plan under which a bankrupt can retain a given fraction of earnings and roll over debt at a lower interest rate than he

[11] Prohibiting saving is meant to capture the seizure of assets in a Chapter 7 bankruptcy. However, this assumption is not quantitatively important, as in our experiments this constraint is binding for only 3 percent of bankrupts. Dropping this restriction on savings has very little effect on the quantitative results, and the fraction of bankrupts who do save is again 3 percent.

[12] Another natural point of comparison—an incomplete market economy where agents cannot default on debts along the lines of that considered in Athreya (2002)—is not feasible in our environment. When expense shocks are sufficiently large, some households will be unable to repay their expense shocks (or save enough a priori). Hence, incorporating expense shocks into an incomplete markets model requires a market arrangement which permits *default along the equilibrium path.*

[13] It should be noted that several European countries changed their laws in the late 1990s.

could access via the market. The household income is subject to the garnishment technology so long as the household remains in bankruptcy.

D. *Timing within the Period*

The timing within a period is as follows. At the beginning of the period, each household realizes its productivity and expense shocks. If the household receives an expense shock κ, the debt of the household is increased (or savings decreased) by κ. The household then decides whether to file for bankruptcy or not. If the household files for bankruptcy, the amount that is garnished is deducted from the earnings, and the consumer is allowed to spend the remainder. Households that declare bankruptcy are unable to save in the period they declared bankruptcy, so they consume all earnings net of garnishment. The new debt level depends on the bankruptcy rule. Households that do not declare bankruptcy choose their net asset holdings for the following period and their current consumption.

E. *Consumer Problem*

We define the consumer's problem recursively. At each date, the household chooses whether to file for bankruptcy, current consumption and next period's debt (savings), taking the bond price schedule as given. In the FS environment, we use three distinct value functions. V is the value of repaying one's debts, while $\bar{V}$ is the value of declaring bankruptcy. We assume that bankruptcy cannot be declared two periods in a row, hence we need one more value function for the period after a bankruptcy.[14] If the household defaults on expense debt, the household's current income is garnished and its debt is rolled over at the fixed interest rate $\bar{r}$: the value of this state is $\bar{W}$.

The value of repaying debts for an age j consumer with debt d and shock realization (z, η, κ) is

$$(3)\quad V_j(d, z, \eta, \kappa) = \max_{c,d'}\left[u\left(\frac{c}{n_j}\right) + \beta E \max\{V_{j+1}(d', z', \eta', \kappa'), \bar{V}_{j+1}(z', \eta')\}\right]$$

$$\text{s.t. } c + d + \kappa \le \bar{e}_j z\eta + q^b(d', z, j)d',$$

where $\bar{V}$ is the value of bankruptcy;

$$(4)\quad \bar{V}_j(z, \eta) = u\left(\frac{c}{n_j}\right) + \beta E \max\{V_{j+1}(0, z', \eta', \kappa'), \bar{W}_{j+1}(z', \eta', \kappa')\},$$

$$\text{where } c = \bar{e}_j z\eta - \Gamma, \quad \Gamma = \gamma\bar{e}_j z\eta,$$

where $\bar{W}$ is the value of not repaying expense shock debt in the period following a bankruptcy;

$$(5)\quad \bar{W}_j(z, \eta, \kappa) = u\left(\frac{c}{n_j}\right) + \beta E \max\{V_{j+1}(d', z', \eta', \kappa'), \bar{V}_{j+1}(z', \eta')\},$$

where

$$c = \bar{e}_j z\eta(1 - \gamma), \quad d' = (\kappa - \gamma\bar{e}_j z\eta)(1 + \bar{r}).$$

When the constraint set in problem (3) is empty, the corresponding value function is equal to $-\infty$.

Let $I_j(d + \kappa, z, \eta)$ denote the decision to declare bankruptcy of an age-j consumer with total debt $d + \kappa$ and current productivity shocks z, η. We assume that borrowers default only if the value of bankruptcy is *strictly* greater than the value of repayment.

The corresponding problem for the NFS environment can be stated with a single value function:

$$(6)\quad V_j^{NFS}(d, z, \eta, \kappa) = \max_{c,d',I}\left[u\left(\frac{c}{n_j}\right) + \beta E V_{j+1}^{NFS}(d', z', \eta', \kappa')\right]$$

$$\text{s.t. } c + d + \kappa \le \bar{e}_j z\eta + q^b(d', z, j)d', \quad \text{if } I = 0;$$

[14] In our parameterization, we assume that each period lasts three years. To capture the U.S. code, we thus have to prohibit bankruptcy in the period immediately following default.

$$c = (1 - \gamma)\bar{e}_j z\eta, \quad \text{if } I = 1;$$

$$d' = \max\{(d + \kappa - \gamma\bar{e}_j z\eta), 0\}(1 + \bar{r}), \quad \text{if } I = 1.$$

F. *Problem of Intermediaries*

Competitive financial markets imply zero expected profits on each loan. Since the law of large numbers holds in our model, ex post realized profits also equal zero. This implies that the price of a bond is determined by the default probability of the issuer and the exogenous risk-free bond price. Let $\theta(d', z, j)$ denote the probability that a household of age j with current productivity shock z and total borrowing d' will declare bankruptcy tomorrow. Without garnishment and with full discharge of debt, the zero profit condition is $q^b(d', z, j) = (1 - \theta(d', z, j))\bar{q}^b$, where $\bar{q}^b = 1/(1 + r^s + \tau)$ is the price of a bond with zero default probability. For positive levels of garnishment, this formula needs to be adjusted for how much lenders can recover from a bankrupt. The *bond price for loans under FS with wage garnishment* is

$$(7) \quad q^b(d', z, j) = (1 - \theta(d', z, j))\bar{q}^b + \theta(d', z, j)E\left(\frac{\Gamma}{d' + \kappa'}\middle| I = 1\right)\bar{q}^b,$$

where $E[(\Gamma/d' + \kappa')|I = 1]$ is the expected rate of recovery through garnishment. We follow the convention that when a household defaults, the amount garnished is allocated proportionately to the repayment of expense debt and personal bonds.

We need to make further adjustments in the NFS case as bankrupts' debts are rolled over. Recall that so long as a household is in arrears, creditors can garnishee a fraction of the earnings in each period. The *bond price under NFS with wage garnishment* is

$$(8) \quad q^{NFS}(d', z, j) = (1 - \theta(d', z, j))\bar{q}^b + \theta(d', z, j) \times E\left(\frac{\Gamma + q(d'', z', j + 1)d''}{d' + \kappa'}\middle| I = 1\right)\bar{q}^b,$$

where $\quad d'' = \max\{d' + \kappa' - \Gamma, 0\}(1 + \bar{r})$.

The key addition to equation (7) is the value of the rolled-over household debt

$$\frac{q(d'', z', j + 1)d''}{d' + \kappa'}.$$

This value is determined by the market value of the rolled-over debt.

G. *Equilibrium*

DEFINITION II.1. *Given a bankruptcy rule and risk-free bond prices* $(q^s, \bar{q}^b)$, *a recursive competitive equilibrium with FS is value functions* V, $\bar{V}$, $\bar{W}$, *policy functions* c, d', $I(d, z, \eta)$, *a default probability* $\theta(d', z, j)$, *and a pricing function* q^b *such that:*

(*a*) *The value functions satisfy the functional equations* (3)–(5), *and* c, d', *and* I *are the associated optimal policy functions.*
(*b*) *The bond prices* q^b *are determined by zero profit condition* (7).
(*c*) *The default probabilities are correct:* $\theta(d', z, j) = E(I_{j+1}(d' + \kappa', z', \eta'))$.

DEFINITION II.2. *A competitive equilibrium with NFS is defined analogously to above, with the modification that* V^{NFS} *has to satisfy the functional equation* (6) *and bond prices* q^{NFS} *are given by equation* (8).

Since the value of declaring bankruptcy (4) is independent of the debt level, and the value of repaying (3) is decreasing in the debt level, the bankruptcy decision in an FS equilibrium follows a simple threshold rule. For every age and income realization, there is a unique level of debt $\bar{d}_j(z, \eta)$ which solves $V_j(\bar{d}, z, \eta, 0) = \bar{V}_j(z, \eta)$. In equilibrium, households repay their debt d if and only if $d + \kappa \leq \bar{d}$.

This makes proving the existence of equilibrium for the Fresh Start environment quite straightforward. Essentially, given any $\bar{q}^b$, there exists a schedule of bond prices q^b such that intermediaries earn zero profits and the solution to consumer's problem is well defined. A formal proof is provided in Livshits, MacGee, and Tertilt (2003).

We compute the equilibrium prices, value, and policy functions by backward induction.

We solve the households' problems given the equilibrium prices which incorporate the default decisions in the following period (starting with the last period of life). We compute the optimal decisions using a grid for the possible asset holdings.

III. Benchmark Parameterization and Results

In this section, we outline the choice of benchmark parameter values. We then compare the benchmark results for the FS system to US data and analyze the basic forces at work in our model.

A. *Benchmark Parameterization*

Households live for 18 periods. Life begins at age 20 and the length of each period is 3 years. The first 15 periods (until age 65) are regular working periods in which people receive income shocks, while the last 3 periods correspond to retirement. We assume that households face no uncertainty during retirement.

The period utility function is $u(c) = (c^{1-\sigma}/1 - \sigma)$, where $1/\sigma$ is the intertemporal elasticity of substitution. We set the annual discount factor equal to 0.94 ($\beta = 0.94^3$) and $\sigma = 2$. The family size life-cycle profile is based on US Census data for 1990. We use the average of several studies of equivalence scales (ES), as reported in Jesús Fernández-Villaverde and Dirk Krueger (forthcoming), to construct an ES life-cycle profile.

The savings interest rate is set equal to 4 percent, which is the average return on capital reported by Ellen R. McGrattan and Edward C. Prescott (2000).[15] This implies a risk-free return on savings for a three-year period of 12.49 percent. The second component of the borrowing interest rate is the transaction cost. We set it equal to 4 percent, which is slightly less than the average cost of making credit card loans reported by David Evans and Richard Schmalensee (1999).[16] This implies a three-year risk-free lending rate of $(1.08)^3 - 1 = 25.97$ percent.

The parameters associated with bankruptcy, γ and $\bar{r}$, also need to be specified. The US bankruptcy codes specify that borrowers must act in "good faith," so that borrowing and immediately filing for bankruptcy is often denied. The parameter γ is intended to capture this fact by requiring that agents repay at least some fraction of their debt. Since we do not have direct data on its magnitude, we calibrate γ so that the debt-income ratio in the benchmark model equals the average ratio of unsecured debt to personal disposable income over the 1995–1999 period, 8.4 percent.[17] This yields a value of 0.355. The annual rollover interest rate, $\bar{r}$, is set to 20 percent.[18]

We parameterize the income process using estimates from the literature. The life-cycle profile of labor income is based on Gourinchas and Parker (2002). A large literature has estimated the volatility of log earnings using the following structure: the log of the persistent idiosyncratic shock follows an AR(1) process

$$\text{(9)} \qquad \ln y_j^i = \ln z_j^i + \ln \eta_j^i + \ln g(X_j^i);$$

$$\ln z_j^i = \rho \ln z_{j-1}^i + \varepsilon_j^i,$$

where $g(X)$ captures the deterministic component of earnings, $\varepsilon_j^i \sim N(0, \sigma_\varepsilon^2)$ and $\eta^i \sim N(0, \sigma_\eta^2)$.[19] We set the benchmark annual value of $\rho = 0.99$, $\sigma_\varepsilon^2 = 0.007$, and $\sigma_\eta^2 = 0.043$. These values are within the range of values reported by Kjetil Storesletten, Chris Telmer, and Amin Yaron (2004), R. G. Hubbard, Skinner, and Stephen P. Zeldes (1994), and Christopher Carroll and Andrew Samwick (1997).

We have to map these annual values into triennial numbers. We set $\rho_3 = 0.99^3$ and $\sigma_{\varepsilon,3}^2 = (1 + \rho^2 + \rho^4) * 0.007 = 0.02$. We discretize the idiosyncratic income shocks using the Tauchen method outlined in Jerome Adda and Russell Cooper (2003). The persistent shock is discretized as a five-state Markov process, with

[15] This value is slightly higher than the average real return on municipal bonds for the U.S., reported by Pierre-Olivier Gourinchas and Jonathan Parker (2002).

[16] This may slightly overestimate the cost, since lending costs are partially offset by fee-charged merchants. This value is comparable to the value used by Steven J. Davis, Felix Kubler, and Paul Willen (forthcoming).

[17] We use revolving credit as reported by the Federal Reserve as our measure of unsecured debt. Recall that this is a model of unsecured consumer debt only.

[18] An interest rate ceiling of 100 percent annually was also imposed on all experiments.

[19] We are abusing notation here, as the variables defined earlier are discrete, whereas here they are continuous.

support $\{z_1, z_2, z_3, z_4, z_5\}$ and age-independent transition matrix $\Pi(z'|z)$. When discretizing the transitory shock, we assume that 10 percent of the population receives a positive (negative) transitory shock each period, and choose the support to match the variance.

We assume that there are no income shocks and no expense shocks during retirement. To capture the dependence of social security on previous income, we assume that retirement earnings are comprised of a lump-sum amount equal to 35 percent of mean earnings in the economy, augmented by 30 percent of a household's earnings in the preretirement period (period 15). These values are within the range of estimates reported by Barbara A. Butrica, Howard M. Iams, and Karen E. Smith (2004).

In our experiments, the expense shocks can take three values: $\kappa \in \{0, \kappa_1, \kappa_2\}$. To calibrate the expense shock, we look at data on out-of-pocket medical bills, divorces, and unplanned (and unwanted) pregnancies. These expenses are both (a) unexpected and (b) frequently cited by bankrupts as the proximate cause of their bankruptcy. While we provide a brief overview of our estimates of the benchmark value of these shocks below, a more detailed discussion of our methodology can be found in Livshits, MacGee, and Tertilt (2003).

We compute the medical expense shock using data on out-of-pocket spending from the 1996 and 1997 waves of the Medical Expenditure Panel Survey (MEPS) and aggregate data from the US Health Care Financing Administration (HCFA).[20] We also take into account unpaid medical bills (using data from the American Hospital Association 1996) by attributing a fraction of the discrepancy between medical charges and expenditures to the expense shock for uninsured individuals. The total medical shock is the sum of 1996 and 1997 and our (bootstrapped) estimate of the medical shock for 1998. These shocks are significant, and a small fraction of households have immense medical bills equal to several times average annual income.

Our estimates of the likelihood of "divorce shocks" and "child shocks" are based on aggregate data of the numbers of households, divorces, and unwanted children from 1996 (US Census Bureau 2000). In calculating the probabilities, we assume these two family events are independent, happen at most once in a three-year period, and that every household is equally likely to be affected. The annual divorce probability is 1.244 percent, which amounts to 3.73 percent per model period.[21] The percentage of U.S. households that is affected by an unplanned and unwanted pregnancy in a given year is 0.5 percent, which amount to a 1.5-percent probability in a model period.[22]

Our estimates of the size of these shocks are also based on aggregate data. The cost of a divorce is based on an average legal fee of \$5,000 and an estimate of the average loss of economies of scale associated with the breakup of a household.[23] We determine a value for a three-year divorce shock of \$36,558. According to the United States Department of Agriculture (1997), the average annual cost of a young child is \$8,000. Assuming that these costs are incurred for three years, we have a shock size of \$24,000.

Since the divorce and child shock amounts are of similar magnitudes, we combine them into one low expense shock by computing the weighted average. We combine this with households that receive a medical shock of equal size, which are 1.874 percent of households. The high expense shock is then based solely on medical bills, and pinned down by the (remaining) right tail of the medical shock distribution.

Table 1 summarizes the expense shocks. The actual values we use are the above values relative to average triennial household disposable income.

[20] Using data from the Health and Retirement Survey (HRS), the Asset and Health Dynamics of the Oldest Old (AHEAD), and tax return data, Feenberg and Skinner (1994) and Eric French and John Jones (forthcoming) also document large out-of-pocket health care expenditures for a small fraction of the population.

[21] This is the number of divorces per 100 households with a head between 23 and 65 years of age, which corresponds to the ages in which shocks are received in the model.

[22] We compute the annual number by multiplying total births per household by the fraction of births that people self-report as unwanted, 0.091 (US Census Bureau 2000).

[23] A typical divorce breaks a three-person household into a one- and a two-person household. Using equivalence scales, this implies an effective income drop of 28 percent (Fernández-Villaverde and Krueger 2002).

TABLE 1—EXPENSE SHOCKS

Shock	Magnitude ($)	Fraction of avg. income	Probability
κ_1	$32,918	0.264	7.104% (π_1)
κ_2	$102,462	0.8218	0.46% (π_2)

B. *Benchmark Results*

In the model, FS is supposed to capture the current US system, while we interpret NFS as a counterfactual experiment—what would happen if the bankruptcy option were taken away. To assess the reasonableness of our benchmark parameter values, we compare two key statistics from the model with the data: the fraction of households declaring bankruptcy and the average borrowing interest rate. All numbers in Table 2 are reported on an annual basis.[24] We also add the debt-income ratio for completeness, but recall that we choose the garnishment parameter to match this fact exactly.

Our model does fairly well in matching the aggregate bankruptcy filings and average borrowing interest rates. The benchmark parameters generate an annual default rate that is 85 percent of that observed in the data (0.71 percent compared to 0.84 percent). However, the data bankruptcy filing rate reported is the number of Chapter 7 nonbusiness filings (reported by the American Bankruptcy Institute) per household, which includes some filings that are primarily due to the failure of nonincorporated small businesses. This is a significant measurement issue; Sullivan, Warren, and Westbrook (1999) find that up to one-fifth of reported nonbusiness filings may be due to the failure of unincorporated small businesses. Hence, the filing rate associated with the benchmark parameters appears to be roughly consistent with the level of US consumer filings. The average borrowing interest rate generated by the model is between the (average over 1995–1999) two rates for unsecured consumer borrowing reported by the Federal Reserve Board: interest rates on two-year person loans and the average interest rate on credit cards carrying balances, both corrected for inflation.

The pattern of defaults in the model is also broadly consistent with US data. Bankrupts in the model have lower earnings than average. The ratio of the mean income of bankrupts to average household income is 0.56, which is similar to the values reported in Sullivan, Warren, and Westbrook (2000). As can be seen from Table 3, most defaults in the model are accompanied by the realization of a negative expense shock. It is worth noting, however, that most households that receive an expense shock do not declare bankruptcy: only one in five of households hit by the small expense shock and roughly 50 percent of households hit by the large shock declare bankruptcy. This suggests that our expense shocks are not so large that they "force" households into bankruptcy, as most households choose to pay the expense shock rather than default.

Our framework generates several interesting life-cycle implications. As can be seen from panel A of Figure 1, our model does a fairly good job of matching bankruptcy rates over the life cycle.[25] Panel B shows consumption and earnings over the life cycle.[26] In our model, consumption tracks earnings fairly closely, as the ability to borrow is limited by the bankruptcy option.

Figure 2 illustrates how the ability to borrow varies over the life cycle. It shows that the maximum amount that people can borrow is hump-shaped over the life cycle. Consumers in our economy are endogenously borrowing-constrained in the sense that even if they increase the face value of the debt, the actual amount received does not increase further. In other words, beyond a certain amount, they face an infinite interest rate.

Since the interest rate varies with the amount borrowed, there is not a single interest rate for each age group. Instead, we calculate average interest rates by age for a *fixed loan size.* (Figure 4B gives such an interest rate life-cycle profile for four different loan amounts.) The u-shaped interest profiles might seem counterintuitive, as

[24] Since each model period corresponds to three years, period default rates from the model are divided by three, while the stock of debt relative to earnings is multiplied by three.

[25] The age-specific filings rates are from Sullivan, Warren, and Westbrook (2000, Table A.4). To facilitate the comparison, we normalize the average filing rate in the data to that of the model.

[26] Data are from Gourinchas and Parker (2002).

TABLE 2—BENCHMARK: MODEL VERSUS DATA

Results	Rule	Debt / Earnings	Defaults	Avg r^b
Benchmark	FS	8.4%	0.71%	11.6%
US data, Avg. 1995–1999	FS	8.4%	0.84%	11.2–12.8

TABLE 3—DEFAULTS BY REASON

	Expense shock			
	Low	High	None	Total
No decrease in income	63.7%	9.9%	1.6%	75.2%
Fall in persistent income only*	8.1%	1.5%	5.3%	14.9%
Negative transitory shock only**	7.0%	1.1%	0.1%	8.3%
Fall in persistent income and negative transitory shock	0.9%	0.2%	0.6%	1.7%
Total	79.7%	12.7%	7.6%	100.0%

* Fall in persistent income = fall in persistent income shock relative to previous period.
** Negative transitory shock = lowest of the three possible values of the transitory income shock.

bankruptcy rates are hump-shaped and one might have expected the interest schedule to follow the same pattern. Note, however, that one reason for the low default rates for older people is the fact that people in these age groups are borrowing very little precisely because borrowing is so expensive. In other words, if old people *were* borrowing as much as people in the middle-age groups, their default rates would be much higher. Middle-aged people, on the other hand, can borrow a given amount at the lowest cost because they have the highest income, and hence repay a bigger fraction of their debt, even if in default.

C. *Comparing FS and NFS*

We now analyze how the current US system (FS) compares to one where people do not have the FS option (NFS). As our measure of welfare, we use the percent increase in lifetime consumption required to equalize expected lifetime utility in both regimes, equivalent consumption variation (ECV). Here a positive number means that FS is the better system. Table 4 shows that in the benchmark economy, welfare is marginally higher under FS than under NFS. This implies that the benefits from increased smoothing across states outweigh the distortion of intertemporal credit markets. The NFS system also generates the expected results for defaults (which are lower) and for debt, which is nearly double the FS level.

Figure 3 nicely illustrates the trade-off between smoothing across time versus states. The life-cycle consumption profile is somewhat steeper under FS, as the bankruptcy option leads to tighter borrowing constraints than under NFS for the average borrower. On the other hand, the bankruptcy option helps people smooth income across states. Except for the young (where borrowing constraints are especially binding) and the retired (who face no uncertainty in the model), the variance in log consumption is much lower in the FS system compared to NFS.

The fact that the endogenous borrowing constraints are much tighter under FS than NFS is clearly visible in Figure 4A. While the borrowing limits for each age are higher for NFS, the gap is largest for younger households. This reflects the fact that NFS allows households to commit to repaying a fraction of their future lifetime income, whereas with FS debt is discharged after one period. Figures 4B and 4C compare the interest rates that agents face in the two systems. Interest rates are higher (for a given amount borrowed) in the FS system, as

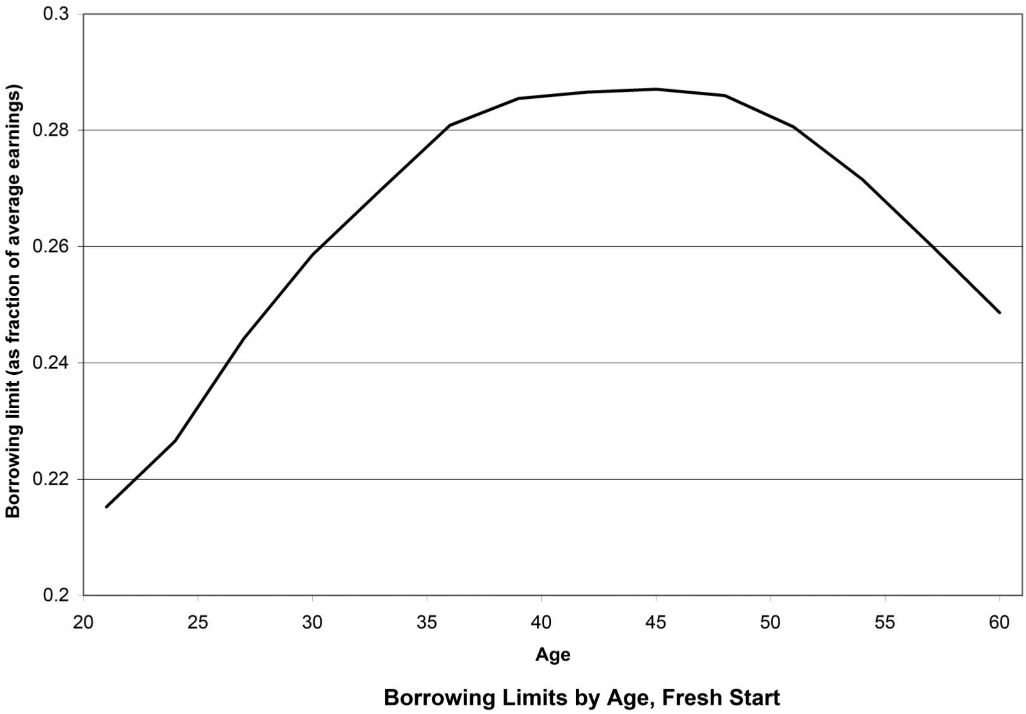

FIGURE 2. BORROWING LIMITS BY AGE, FRESH START

TABLE 4—BENCHMARK: FS VERSUS NFS

Results	Rule	Debt / Earnings	Defaults	Better rule	ECV
Benchmark	FS	8.4%	0.71%	FS	0.06%
	NFS	14.8%	0.53%		

more people default on their debt. This difference is largest for young agents, and is driven by the greater ability of young households to commit to repaying out of future income under NFS.

IV. Importance of Uncertainty and Life Cycle

A key insight of this paper is that the evaluation of bankruptcy regimes is sensitive both to the nature and the magnitude of the idiosyncratic uncertainty and to life-cycle considerations. We now summarize several experiments that illustrate this point.

A. *Expense Uncertainty*

We find that expense uncertainty plays a crucial role in evaluating alternative bankruptcy regimes. When we set expense shocks to zero, life-long liability for debt (NFS) is preferable to permitting the discharge of debt (FS). If expense shocks are sufficiently high, however, FS is preferable to NFS. In Figure 5, we display two sets of experiments: scaling the magnitudes of both shocks (κ_1, κ_2) and scaling the probabilities (π_1, π_2) by a factor. FS becomes relatively more attractive as the size of the shock increases (and the attractiveness of FS

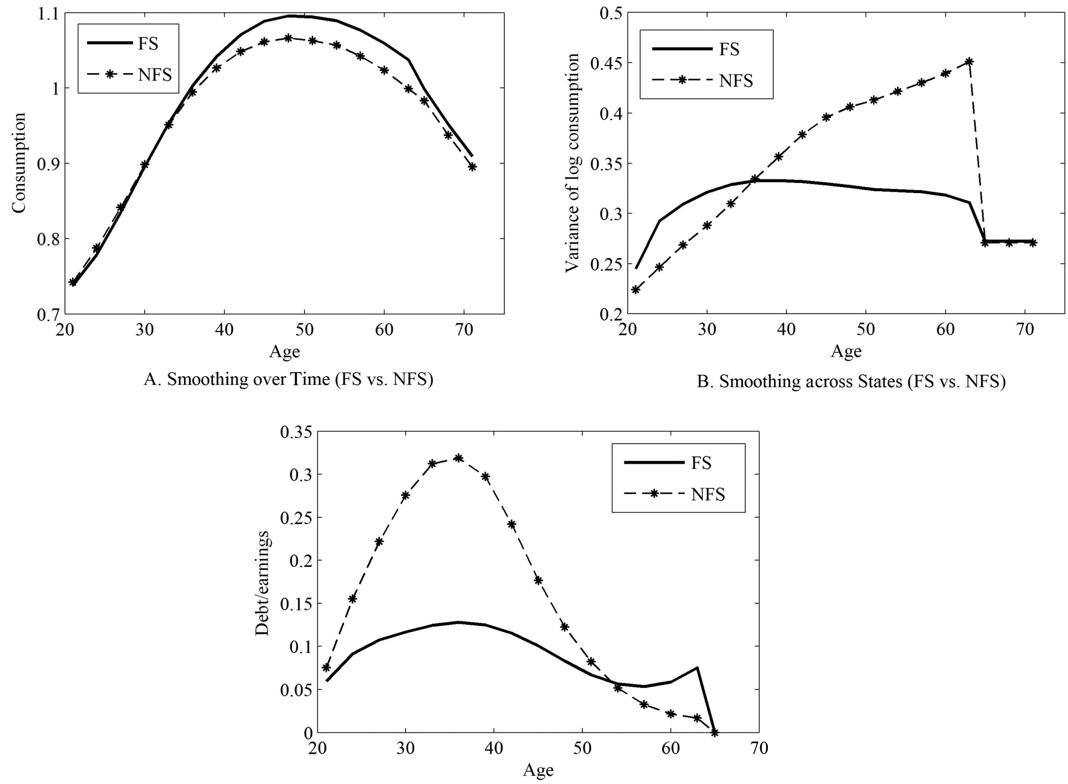

A. Smoothing over Time (FS vs. NFS)

B. Smoothing across States (FS vs. NFS)

C. Debt/Earnings over the Life Cycle (FS vs. NFS)

FIGURE 3

increases faster in the size of the shock than in the probabilities).

We also find that how we model expense shocks matters. One could argue that expense shocks such as divorce should be proportional to the persistent component of earnings. This concern may be important, as we find that if both expense shocks are made proportional (keeping the probabilities and the average size of the shock unchanged), then NFS is better than FS by 0.61 percent of ECV. The MEPS data, however, suggest that out-of-pocket expenditures are highest for the lowest and the highest earning quintiles. With this in mind, we ran an experiment in which only the small expense shock was proportional to earnings, and found that the advantage of NFS relative to FS came down to 0.41 percent of ECV.

These findings have two important implications. First, they suggest that it is not implausible that the debt discharge provision in the current US bankruptcy law may be welfare improving. This conclusion differs from that of Athreya (2002) (and others) who abstract from expense uncertainty and find that eliminating bankruptcy in the United States would increase welfare.[27] Moreover, our findings lend support to the views advanced by sociologists and lawyers such as Sullivan, Warren, and Westbrook (2000) that bankruptcy plays an important role in providing a safety net against bad luck for Americans.

B. *Earnings Uncertainty*

Transitory Shocks.—Transitory shocks to earnings have little effect on lifetime wealth and can be smoothed over time. Intuition suggests

[27] It should be noted that Athreya (2002) also differs from our model in the pricing of debt.

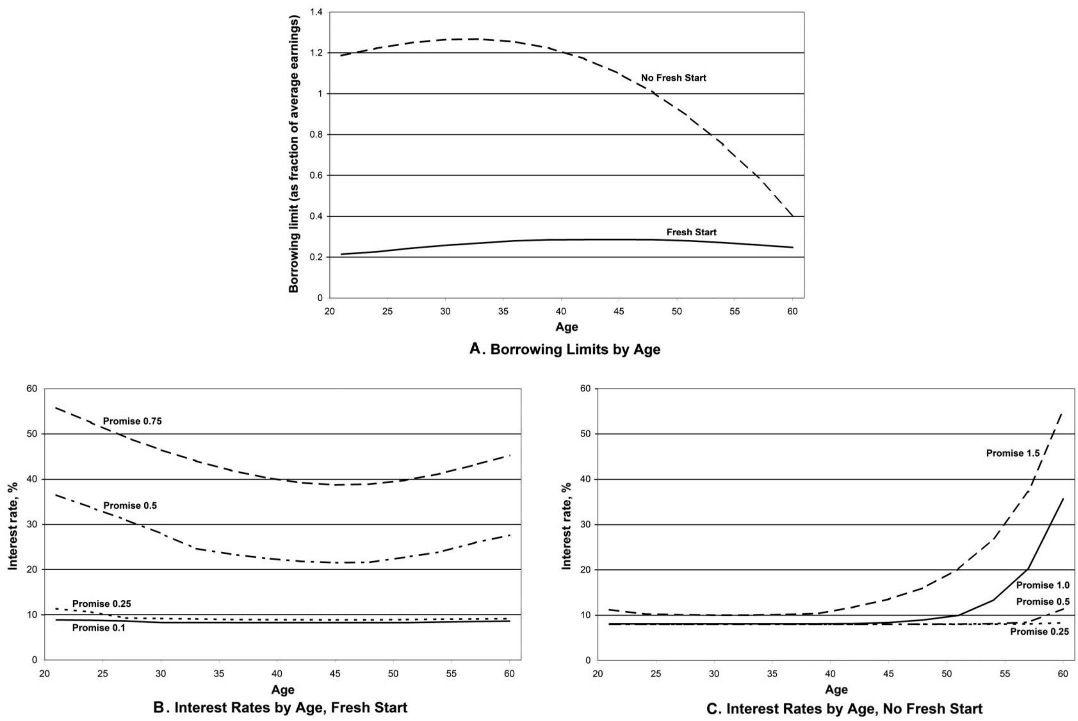

FIGURE 4

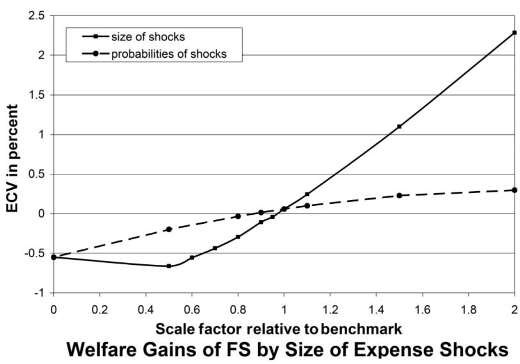

FIGURE 5. WELFARE GAINS OF FS BY SIZE OF EXPENSE SHOCKS

that this should make the ability to smooth intertemporally relatively more important than having an option to walk away from debt. As a result, we would expect that an increase in the variance of temporary shocks would make NFS more attractive relative to FS.

This logic is reflected in our experiments as reported in Table 5. Increasing the variance of the transitory shock makes NFS more attractive relative to FS. The other variables move as one would expect. Under FS, borrowing at first rises as the variance of transitory shocks increases. For variances above the benchmark, however, borrowing declines as households seek to self-insure against the increase in uncertainty. In contrast, borrowing rises monotonically under NFS with the variance in transitory shocks. Bankruptcy filings and the average borrowing interest rate under FS also experience a small rise with the variance of transitory shocks.

Persistent Shocks.—The effect of changes in the variance on the FS versus NFS comparison is much more nuanced for persistent income shocks, and depends critically upon the persistence of the income process. Increases in the variance of persistent income shock make FS less attractive to lower-income households, while higher-income households' preference for FS increases. However, both the magnitude of these effects and the cut-off group vary systematically with the level of persistence. As a result, the impact of changes

TABLE 5—EFFECTS OF THE TRANSITORY SHOCK

Var. (σ_η^2)	0	0.01	0.02	0.03	0.043	0.05	0.08	0.15
ECV, %	0.53	0.42	0.32	0.22	0.06	−0.01	−0.36	−1.14

TABLE 6—EFFECTS OF THE PERSISTENT SHOCK: ECV FS OVER NFS, %

	ρ		
$\sigma^2_{\varepsilon,3}$	0.98	0.99	0.995
0.005	−0.167	−0.086	0.108
0.01	−0.196	−0.027	0.387
0.02	−0.285	0.058	1.029
0.04	−0.549	0.123	1.969

in the variance on the FS-NFS comparison changes with the degree of persistence of the income process.

This relationship is illustrated by the experiments reported in Table 6. The middle column of Table 6 reports the effects of changes in the variance for the benchmark level of persistence, $\rho = 0.99$. As the variance increases, FS becomes better compared to NFS. For higher persistence, a similar relationship is observed (see the last column). However, for lower persistence (e.g. $\rho = 0.98$), increased variance works in favor of NFS.

To understand these results, one has to look at the trade-off between smoothing over time and smoothing across states for different (persistent) income groups. The lowest (persistent) income group's preference for NFS is increasing in the variance. As the variance increases, the expected value of lifetime income increases relative to current income. This generates an increased desire for lower-income households to borrow, which works in favor of NFS. The lower the persistence of the income process, the larger the magnitude of this force. In contrast, the highest-income groups have a precautionary saving motive which is increasing in the variance of income. This precautionary saving motive makes the tighter borrowing constraints associated with FS less costly. The middle income group's preference is a mix of these two forces. In our experiments, at high levels of persistence the insurance motive becomes more valuable for these groups relative to intertemporal smoothing.

C. *Importance of the Life Cycle*

Finally, we provide two experiments, reported in Table 7, that demonstrate that the life-cycle dimension is very important. This shows that a welfare comparison based on a model without a life-cycle component needs to be interpreted with caution.

First, suppose that family size did not vary over the life cycle. This makes FS less attractive relative to NFS, because ignoring the hump-shaped family size profile increases the desire of the young to borrow against future high earnings. Note that while aggregate borrowing does increase under FS compared to the benchmark, it goes up by (proportionally) more in the NFS economy.

Suppose instead that the labor income profile were flat, i.e., there were no life-cycle aspect to earnings. This makes FS much more attractive relative to NFS. The logic is the same: with a flat income profile, there is no need to smooth over time. This is reflected in a much lower level of debt relative to the benchmark economy and a reduction in the ratio of FS to NFS debt.

V. Conclusion

In this paper, we develop a formal model of consumer bankruptcy with a competitive lending market for unsecured credit, and use it to quantitatively analyze different consumer bankruptcy rules. Our model generates interest rates that differ across types of consumers and also depend on a consumer's total debt. For reasonable parameter values, we match the level of unsecured consumer debt and bankruptcy filings rates in the United Sates fairly well.

There are two key messages with regards to the evaluation of bankruptcy regimes. The first is that the welfare comparison of bankruptcy regimes varies with both the nature and extent of uncertainty that households face over the life cycle. We find that incorporating expense shocks lends some support to the view that the FS provisions of the US bankruptcy act are welfare improving. We

TABLE 7—IMPORTANCE OF LIFE-CYCLE PROFILE

		Debt/GDP	Bankruptcies	ECV
Flat family profile	FS	8.8%	0.71%	−0.24%
	NFS	16.17%	0.51%	
Flat earnings profile	FS	5.2%	0.56%	+0.63%
	NFS	7.0%	0.46%	

also find that the life cycle plays a key role in our analysis. We feel that further work on both improving our measurements of uncertainty and on the details of the modeling of the bankruptcy rules might be needed before one can take a stand on the recent changes to US bankruptcy legislation.

It is also worth noting that there are several aspects we have abstracted from that deserve attention in future work. One is the extent to which bankruptcy rules affect households' decisions to insure against (or to take actions to mitigate against) shocks. This might imply that the shock process itself is endogenous to the system. A second is the effect of durable assets (such as houses and cars) on the market for unsecured credit. While most Chapter 7 bankrupts own no exempt assets, the existence of secured assets may still affect a household's ability to smooth consumption in response to transitory shocks to income and wealth. This is a topic we hope to explore in future work.

REFERENCES

Adda, Jérôme, and Russell Cooper. 2003. *Dynamic Economics: Quantitative Methods and Applications*. Cambridge, MA and London: MIT Press.

Aiyagari, S. Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Quarterly Journal of Economics*, 109(3): 659–84.

Alexopoulos, Michelle, and Ian Domowitz. 1998. "Personal Liabilities and Bankruptcy Reform: An International Perspective." *International Finance Journal*, 1(1): 127–59.

American Bankruptcy Institute. 2005. "Major Consumer Bankruptcy Effects of the 2005 Reform Legislation." http://www.Abiworld.Org/Pdfs/S256/Mainpoints11.pdf (accessed August 22, 2006).

Athreya, Kartik B. 2002. "Welfare Implications of the Bankruptcy Reform Act of 1999." *Journal of Monetary Economics*, 49(8): 1567–95.

Attanasio, Orazio P., and Guglielmo Weber. 1995. "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey." *Journal of Political Economy*, 103(6): 1121–57.

Butrica, Barbara A., Howard M. Iams, and Karen E. Smith. 2003. "The Changing Impact of Social Security on Retirement Income in the United States." *Social Security Bulletin*, 65(3): 1–13.

Carroll, Christopher D., and Andrew A. Samwick. 1997. "The Nature of Precautionary Wealth." *Journal of Monetary Economics*, 40(1): 41–71.

Chatterjee, Satyajit, Dean Corbae, Makoto Nakajima, and José-Víctor Ríos-Rull. 2005. "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default." Philadelphia Federal Reserve Bank Working Paper 05-18.

Davis, Steven J., Felix Kubler, and Paul Willen. 2006. "Borrowing Costs and the Demand for Equity over the Life Cycle." *Review of Economics and Statistics*, 88(2): 348–62.

Domowitz, Ian, and Robert L. Sartain. 1999. "Determinants of the Consumer Bankruptcy Decision." *Journal of Finance*, 54(1): 403–20.

Dubey, Pradeep, John Geanakoplos, and Martin Shubik. 2005. "Default and Punishment in General Equilibrium." *Econometrica*, 73(1): 1–37.

Evans, David S., and Richard Schmalensee. 1999. *Paying with Plastic: The Digital Revolution in Buying and Borrowing*. Cambridge, MA and London: MIT Press.

Fay, Scott, Erik Hurst, and Michelle J. White. 2002. "The Household Bankruptcy Decision." *American Economic Review*, 92(3): 706–18.

Feenberg, Daniel, and Jonathan Skinner. 1994. "The Risk and Duration of Catastrophic Health Care Expenditures." *Review of Economics and Statistics*, 76(4): 633–47.

Fernández-Villaverde, Jesús, and Dirk Krueger. Forthcoming. "Consumption over the Life Cycle: Some Facts from Consumer Expenditure Survey Data." *Review of Economics and Statistics*.

French, Eric, and John Bailey Jones. 2004. "On the Distribution and Dynamics of Health Care Costs." *Journal of Applied Econometrics*, 19(6): 705–21.

Gourinchas, Pierre-Oliver, and Jonathan A. Parker. 2002. "Consumption over the Life Cycle." *Econometrica*, 70(1): 47–89.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. 1994. "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving." *Carnegie-Rochester Conference Series on Public Policy*, 40: 59–125.

Kehoe, Timothy J., and David K. Levine. 1993. "Debt-Constrained Asset Markets." *Review of Economic Studies*, 60(4): 865–88.

Kocherlakota, Narayana R. 1996. "Implications of Efficient Risk Sharing without Commitment." *Review of Economic Studies*, 63(4): 595–609.

Li, Wenli, and Pierre-Daniel Sarte. 2006. "U.S. Consumer Bankruptcy Choice: The Importance of General Equilibrium Effects." *Journal of Monetary Economics*, 53(3): 613–31.

Livshits, Igor, **James MacGee, and Michèle Tertilt**. 2003. "Consumer Bankruptcy: A Fresh Start." Federal Reserve Bank of Minneapolis Working Paper 617.

Luckett, Charles A. 2002. "Personal Bankruptcies." In *Impact of Public Policy on Consumer Credit*, ed. Thomas Durkin and Michael Staten, 69–102. Norwell and Dordrecht: Kluwer Academic Publishers.

McGrattan, Ellen R., and Edward C. Prescott. 2000. "Is the Stock Market Overvalued?" *Federal Reserve Bank of Minneapolis Quarterly Review*, 24(4): 20–40.

National Bankruptcy Review Commission. 1997. "Bankruptcy: The Next Twenty Years. Final Report of the Bankruptcy Review Commission." Technical Report. http://govinfo.library.unt.edu/nbrc/reporttitlepg.html.

Niemi-Kiesilainen, Johanna. 1997. "Changing Directions in Consumer Bankruptcy Law and Practice in Europe and USA." *Journal of Consumer Policy*, 20(2): 133–42.

Pavan, Marina. 2005. "Consumer Durables and Risky Borrowing: The Effects of Bankruptcy Protection." Boston College Working Paper in Economics 573.

Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron. 2004. "Consumption and Risk Sharing over the Life Cycle." *Journal of Monetary Economics*, 51(3): 609–33.

Sullivan, Teresa A., Elizabeth Warren, and Jay Lawrence Westbrook. 1999. *As We Forgive Our Debtors: Bankruptcy and Consumer Credit in America*. Washington, DC: Beard Books.

Sullivan, Teresa A., Elizabeth Warren, and Jay Lawrence Westbrook. 2000. *The Fragile Middle Class: Americans in Debt*. New Haven and London: Yale University Press.

U.S. Census Bureau. 2000. *Statistical Abstract of the United States* http://www.census.gov/prod/2001pubs/statab/sec02.pdf.

U.S. Department of Agriculture. 1997. *Expenditures on Children by Families, 1996 Annual Report*. Miscellaneous Publication No. 1528–1996. http://www.cnpp.usda.gov/Publications/CRC/crc1996.pdf.

Warren, Elizabeth, Teresa A. Sullivan, and Melissa B. Jacoby. "Medical Problems and Bankruptcy Filings." *Norton Bankruptcy Law Advisor May 2000.* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=224581.

Zame, William R. 1993. "Efficiency and the Role of Default When Security Markets Are Incomplete." *American Economic Review*, 83(5): 1142–64.